

Adaptive Video Streaming in MU-MIMO Networks

D. Bethanabhotla, G. Caire and M. J. Neely

Abstract—We consider extensions and improvements on our previous work [1] on dynamic adaptive video streaming in a multi-cell multiuser “small cell” wireless network. In [1] we treated the case of single-antenna base stations and, starting from a network utility maximization (NUM) formulation, we devised a “push” scheduling policy, where users place requests to sequential video chunks to possibly different base stations with adaptive video quality, and base stations schedule their downlink transmissions in order to stabilize their transmission queues. In this paper we consider a “pull” strategy, where every user maintains a request queue, such that users keep track of the video chunks that are effectively delivered. The pull scheme allows to download the chunks in the playback order without skipping or missing them. In addition, motivated by the recent/forthcoming progress in small cell networks (e.g., in wave-2 of the recent IEEE 802.11ac standard), we extend our dynamic streaming approach to the case of base stations capable of multiuser MIMO downlink, i.e., serving multiple users on the same time-frequency slot by spatial multiplexing. By exploiting the “channel hardening” effect of high dimensional MIMO channels, we devise a low complexity user selection scheme to solve the underlying max-weighted rate scheduling, which can be easily implemented and runs independently at each base station. Through simulations, we show MIMO gains in terms of video streaming QoE metrics like the pre-buffering and re-buffering times.

I. INTRODUCTION

Wireless data traffic is predicted to increase dramatically in the next few years, up to two orders of magnitude by 2020 [2]. This increase is mainly due to on-demand video streaming, enabled by multimedia devices such as tablets and smartphones. It is well understood that the current trend of cellular technology (e.g., LTE [3]) cannot cope with such traffic increase, unless the density of the deployed wireless infrastructure is increased correspondingly. This motivates the recent flurry of research on massive and dense deployment of base station antennas, either in the form of “massive MIMO” solutions (hundreds of antennas at each cell site [4]) or in the form of very dense small-cell networks [5]. While discussing the relative merits of these approaches is out of the scope of this paper, we mention here that the ‘small-cell MU-MIMO’ solution seems to be most immediately applicable, since it can leverage the rapidly evolving MIMO technology advances of wireless local-area networks standards, such as IEEE 802.11 ac with MU-MIMO capability [6]. This paper focuses on the problem of dynamic adaptive video streaming in a wireless network formed by a number of small-cell base station (helper) nodes, employing MU-MIMO technology, serving multiple wireless users over a given geographic coverage area and on the same shared channel bandwidth.

Contributions: We introduce the notion of a *request queue*. This is a virtual queue, maintained by each user, that serves to sequentially request video chunks from helper nodes, such that the choice of the helper node and the quality at which

each video chunk is requested can be adaptively adjusted. In order to obtain a dynamic policy with provable optimality properties, we formulate a NUM problem [7] and solve it using the drift plus penalty (DPP) approach in the framework of Lyapunov Optimization [8]. The obtained solution is provably asymptotically optimal in a per-sample path sense (i.e., without assuming stationarity and ergodicity of the underlying network state process [8]). Furthermore, it naturally decomposes into two sub-policies: “congestion control” which is implemented independently at every user and, “transmission scheduling” which is implemented independently at every helper. The congestion control decision consists of each user adaptively selecting the video quality level of the chunks and “virtually” placing them in its request queue. Note that this does not mean the user already has the chunk, but the chunk is “virtually” placed in the request queue and will be taken out when it is effectively delivered to the user. In addition, the user broadcasts its request queue length to the helpers in its current vicinity and requests from them only those bits which are at the ‘head of line’ of its queue. In this way, the user always downloads chunks in the playback order and does not skip any of them. The transmission scheduling decision consists of each “MU-MIMO” base station greedily choosing the subset of users to be served by multiuser MIMO spatial multiplexing based on the request queue length information broadcasted by the users in its vicinity.

II. SYSTEM MODEL

We consider a discrete, time-slotted wireless network with multiple users and multiple helper stations sharing the same bandwidth. The network is defined by a bipartite graph $\mathcal{G} = (\mathcal{U}, \mathcal{H}, \mathcal{E})$, where \mathcal{U} denotes the set of users, \mathcal{H} denotes the set of helpers, and \mathcal{E} contains edges for all pairs (h, u) such that there exists a potential transmission link between $h \in \mathcal{H}$ and $u \in \mathcal{U}$. We denote by $\mathcal{N}(u) \subseteq \mathcal{H}$ the neighborhood of user u , i.e., $\mathcal{N}(u) = \{h \in \mathcal{H} : (h, u) \in \mathcal{E}\}$. Similarly, $\mathcal{N}(h) = \{u \in \mathcal{U} : (h, u) \in \mathcal{E}\}$. Each user $u \in \mathcal{U}$ requests a video file f_u from a library \mathcal{F} of possible files. Each video file is formed by a sequence of chunks. Each chunk corresponds to a group of pictures (GOP) that are encoded and decoded as stand-alone units [9]. Chunks have a fixed playback duration, given by $T_{\text{gop}} = (\# \text{ frames per GOP})/\eta$, where η is the frame rate, expressed in frames per second. The streaming process consists of transferring chunks from the helpers to the requesting users such that the playback buffer at each user contains the required chunks at the beginning of each chunk playback deadline. The playback starts after a short pre-buffering time, during which the playback buffer is filled by a determined amount of ordered chunks. The details relative to pre-buffering and chunk playback deadlines are discussed in Section V.

Each file $f \in \mathcal{F}$ is encoded at a finite number of different quality levels $m \in \{1, \dots, N_f\}$. This is similar to the implementation of several current video streaming technologies, such as Microsoft Smooth Streaming and Apple HTTP Live Streaming [10]. Due to the variable bit rate (VBR) nature of video coding [11], the quality-rate profile of a given file f may vary from chunk to chunk. We let $D_f(m, t)$ and $B_f(m, t)$ denote the video quality measure (e.g., see [12]) and the number of bits per pixel for file f at chunk time t and quality level m respectively. Letting N_{pix} denote the number of pixels per frame, a chunk contains $k = \eta T_{\text{gop}} N_{\text{pix}}$ pixels. Hence, the number of bits in the t -th chunk of file f , encoded at quality level m , is given by $kB_f(m, t)$.

Each user maintains a *request queue* Q_u of bits that it wants to download possibly from different helpers it might associate with during its video streaming session. Note that Q_u is different from the actual playback buffer Ψ_t (see Section V) and acts like a bookkeeper by maintaining the list of chunks that have been requested but not downloaded yet. In addition, each user u independently makes the decision of choosing the quality mode $m_u(t)$ for chunk time t . This choice affects the choice of the quality $D_{f_u}(m_u(t), t)$ and the size $kB_{f_u}(m_u(t), t)$ of the chunk t that it places in its request queue Q_u . The dynamics of Q_u for each user u is given by:

$$Q_u(t+1) = \max\{Q_u(t) - n\mu_u(t) + kB_u(t), 0\} \quad \forall u \in \mathcal{U}, \quad (1)$$

where n denotes the number of physical layer channel symbols corresponding to a time slot of duration T_{gop} , and $n\mu_u(t) = \sum_{h \in \mathcal{N}(u)} n\mu_{hu}(t)$ is the aggregate number of video-encoded bits per time slot that the user is able to download from its neighboring helpers. Note that $\mu_{hu}(t)$ is the channel coding rate (bits/channel symbol) of the transmission from helper h to user u at time t . Here, we assume that user u at time t can receive $\mu_u(t) = \sum_{h \in \mathcal{N}(u)} \mu_{hu}(t)$ bits/channel symbol by simultaneously downloading $n\mu_{hu}(t)$ bits from helpers h in $\mathcal{N}(u)$. Although this is not implemented in 802.11 networks, receiving multiple data streams from multiple base station is definitely possible (e.g., in CDMA system with macro diversity). In any case, in this paper, we do not assume that the user is able to perform joint or successive decoding of the multiple streams: when a user is served by multiple helpers on the same time slot, each stream is decoded independently and treats everything else as Gaussian noise. Note that the $kB_u(t)$ bits that user u places in its request queue Q_u correspond to the chunk t while the $n\mu_u(t)$ bits that it downloads from its neighboring helpers correspond to the chunks at the ‘head of the line’ of Q_u . In this way, the user can download chunk $t+1$ only after having downloaded chunk t and consequently receive chunks in order of playback. This is in contrast to the ‘push’ scheme in [1] where a user, depending on the queue lengths of the neighboring helpers, may download chunks out of order. We consider two possible physical layer systems in this paper. PHY A is the system which was already considered in [1], where each helper h has a single antenna and serves its neighboring users $u \in \mathcal{N}(h)$ using orthogonal FDMA/TDMA. System PHY B is a significant extension of PHY A where now each helper implements MU-MIMO to serve its neighboring

users and time shares among different possible subsets of served users for zero-forcing beamforming. We describe only PHY B and refer the reader to [1] for a description of PHY A (which will be used as a term of comparison to show the improvement due to MU-MIMO downlink).

PHY B : MU-MIMO Base Stations

In this system, each helper h , with a large number of antennas M installed, implements MU-MIMO to serve the users associated to it, i.e., $\mathcal{N}(h)$. As a result, helper h can serve simultaneously, in the spatial domain, any subset of size not larger than $\min\{M, |\mathcal{N}(h)|\}$ of the clients in $\mathcal{N}(h)$. We further assume that each base station performs linear zero-forcing beamforming (LZFBF) to the set of selected users (referred to in the following as “active users”) it serves simultaneously. In addition, each helper h operates in TDMA over all possible active user subsets $\mathbf{S}_h \subseteq \mathcal{N}(h)$. We use $\mathbf{S}_h(\tau)$ to denote the subset that is chosen in OFDM resource block τ , which is different from chunk time t and in fact much smaller than t . This assumption is motivated by realistic typical system parameters, where the time and frequency selective wireless channel fading coherence time \times bandwidth product is small with respect to the number of signal dimensions spanned by the transmission of a video chunk. We let $S_h(\tau)$ to denote the cardinality of $\mathbf{S}_h(\tau)$. Under the assumptions that $M, S_h(\tau) \rightarrow \infty$ with a fixed ratio $\frac{S_h(\tau)}{M} < 1$, random matrix theory results can be invoked to show that for a given choice of subset $\mathbf{S}_h(\tau)$ and under the reasonable assumption that the power P_h is equally shared across the user streams in $\mathbf{S}_h(\tau)$, the vector of rates achieved by all the users in $\mathcal{N}(h)$ is given by

$$C_{hu}(\mathbf{S}_h(\tau), t) = \begin{cases} 0 & \text{if } u \notin \mathbf{S}_h(\tau) \\ \log \left(1 + \frac{g_{hu}(t)(M - S_h(\tau) + 1)P_h}{S_h(\tau)(1 + \sum_{h' \neq h} P_{h'} g_{h'u}(t))} \right) & \text{if } u \in \mathbf{S}_h(\tau) \end{cases} \quad (2)$$

In fact, it is known that the asymptotics kick in very quickly making the rates in (2) achievable for practical values of M and $S_h(\tau)$. Notice that the rate expression is independent of the small scale fading coefficients which vary at a faster time scale (i.e. at the scale of OFDM resource blocks). This is because of using a large number of antennas M at the helpers which renders a large $M \times S_h(\tau)$ random channel matrix \mathbf{H} of i.i.d complex Gaussian distributed small scale fading coefficients in every OFDM resource block τ . When each helper performs LZFBF in every resource block τ , it turns out that rate expressions involve calculating the reciprocals of the diagonal elements of the inverse Wishart matrix $(\mathbf{H}^H \mathbf{H})^{-1}$. Due to the large size of the matrix \mathbf{H} and under the assumption $\frac{S_h(\tau)}{M} < 1$, random matrix theory results can be invoked to show that the diagonal elements ‘harden’ at a deterministic value. This results in deterministic rate expressions as in (2) which are independent of \mathbf{H} and are just dependent on the large scale path loss coefficients $g_{hu}(t)$ (see [13] for more details).

Let $\mathbf{C}_h(\mathbf{S}_h(\tau), t)$ be the $|\mathcal{N}(h)|$ dimensional vector whose elements are the rates $C_{hu}(\mathbf{S}_h(\tau), t)$ as given in (2) achieved

by every user $u \in \mathcal{N}(h)$ when the helper h beamforms to the active user subset $\mathbf{S}_h(\tau)$ in resource block τ . In addition, let $\boldsymbol{\mu}_h(t)$ be the vector obtained by averaging the rate vector $\mathbf{C}_h(\mathbf{S}_h(\tau), t)$ scheduled by helper h to users in $\mathcal{N}(h)$ over all the resource blocks τ in chunk time t . Since we assume that helper h serves its neighboring users $u \in \mathcal{N}(h)$ by sharing the resource blocks over all possible active user subsets of $\mathcal{N}(h)$, $\boldsymbol{\mu}_h(t)$ is constrained to lie in the LZFBF-achievable region of the underlying MIMO broadcast channel of base station h and users $\mathcal{N}(h)$. This yields the transmission rate constraint

$$\boldsymbol{\mu}_h(t) \in \text{coh}\{\mathbf{C}_h(\mathbf{S}_h, t) : \mathbf{S}_h \subseteq \mathcal{N}(h)\} \quad \forall h \in \mathcal{H} \quad (3)$$

where coh is the short-hand notation for ‘convex hull’. Note that the above region is a convex polytope in $|\mathcal{N}(h)|$ dimensions obtained by taking the convex hull of $2^{|\mathcal{N}(h)|} - 1$ rate vectors $\mathbf{C}_h(\mathbf{S}_h, t)$.

In both systems PHY A and PHY B, the slow fading gain $g_{hu}(t)$ models path loss and shadowing between helper h and user u , and it is assumed to change slowly in time. For a scenario typical of small cell networks, where users are nomadic (e.g., moving at walking speed), the slow fading coefficients change on a time-scale of the order of 10s (i.e., ≈ 20 scheduling slots assuming a realistic and common value of $T_{\text{gop}} = 0.5\text{s}$). This time scale is much slower than the coherence of the small-scale fading, but is comparable with the duration of the video chunks. Therefore, variations of these coefficients during a video streaming session (e.g., due to user mobility) are relevant.

We let $\boldsymbol{\omega}(t)$ denote the network state at time t , defined as $\boldsymbol{\omega}(t) = \{g_{hu}(t), D_{fu}(\cdot, t), B_{fu}(\cdot, t) : \forall (h, u) \in \mathcal{E}\}$. Let $A_{\boldsymbol{\omega}(t)}$ be the set of feasible control actions, dependent on the current network state $\boldsymbol{\omega}(t)$, and let $\alpha(t) \in A_{\boldsymbol{\omega}(t)}$ be a control action, comprising the vector $\boldsymbol{\mu}(t)$ with elements $\mu_u(t)$, the quality modes $m_u(t) \forall u \in \mathcal{U}$ and the vector $\mathbf{B}(t)$ with elements $B_{fu}(m_u(t), t)$. A control policy for the system at hand is a sequence of control actions $\{\alpha(t)\}_{t=0}^{\infty}$ where at each time t , $\alpha(t) \in A_{\boldsymbol{\omega}(t)}$.

III. DYNAMIC STREAMING POLICY DESIGN

In the proposed NUM problem, the goal consists of designing a control policy which maximizes a concave utility function of the time averaged *video qualities* of the users, subject to keeping the request queues at every user stable. Using the notation $\bar{x} := \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}[x(\tau)]$ for the long-term time average expectation of a quantity x , Define $\bar{D}_u := \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}[D_{fu}(m_u(\tau), \tau)]$ and $\bar{Q}_u := \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}[Q_u(\tau)]$ as the long-term time average of the expected quality level and the expected request queue length respectively at user u . Let $\phi_u(\cdot)$ be a concave, continuous, and non-decreasing function defining utility vs. video quality for user $u \in \mathcal{U}$. The goal is to solve:

$$\text{maximize} \quad \sum_{u \in \mathcal{U}} \phi_u(\bar{D}_u) \quad (4)$$

$$\text{subject to} \quad \bar{Q}_u < \infty \quad \forall u \in \mathcal{U} \quad (5)$$

$$\alpha(t) \in A_{\boldsymbol{\omega}(t)} \quad \forall t, \quad (6)$$

where constraint (5) corresponds to the *strong stability* condition for all the queues Q_u which ensures that all the requested chunks will be eventually delivered. Problem (4) – (6) can be solved using the stochastic optimization theory of [8]. Since it involves maximizing a *function* of time averages, it is convenient to transform it into an equivalent problem that involves maximizing a single time average instead of a function of time averages. Then, the *drift plus penalty* framework of [8] can be applied. This transformation is achieved through the use of auxiliary variables $\gamma_u(t)$ and corresponding virtual queues $\Theta_u(t)$ with buffer evolution:

$$\Theta_u(t+1) = \max\{\Theta_u(t) + \gamma_u(t) - D_{fu}(m_u(t), t), 0\}. \quad (7)$$

Consider the transformed problem:

$$\text{maximize} \quad \sum_{u \in \mathcal{U}} \overline{\phi_u(\gamma_u)} \quad (8)$$

$$\text{subject to} \quad \bar{Q}_u < \infty \quad \forall u \in \mathcal{U} \quad (9)$$

$$\bar{\gamma}_u \leq \bar{D}_u \quad \forall u \in \mathcal{U} \quad (10)$$

$$D_u^{\min} \leq \gamma_u(t) \leq D_u^{\max} \quad \forall u \in \mathcal{U} \quad (11)$$

$$\alpha(t) \in A_{\boldsymbol{\omega}(t)} \quad \forall t \quad (12)$$

where D_u^{\max} is a uniform upper bound on the maximum quality $D_{fu}(N_{fu}, t)$ and D_u^{\min} is a lower bound on the minimum quality $D_{fu}(1, t)$, for all chunk times t . Notice that constraints (10) correspond to stability of the virtual queues Θ_u , since $\bar{\gamma}_u$ and \bar{D}_u are the time-averaged arrival rate and the time-averaged service rate for the virtual queue given in (7).

Let $\mathbf{Q}(t)$, $\boldsymbol{\Theta}(t)$, $\boldsymbol{\gamma}(t)$ and $\mathbf{D}(t)$ denote the column vectors with the elements $Q_u(t)$, $\Theta_u(t)$, $\gamma_u(t)$ and $D_{fu}(m_u(t), t)$ respectively. Let $\mathbf{G}(t) = [\mathbf{Q}^T(t), \boldsymbol{\Theta}^T(t)]^T$ be the combined vector of queue backlog vectors and define the quadratic Lyapunov function $L(\mathbf{G}(t)) := \frac{1}{2} \mathbf{G}^T(t) \mathbf{G}(t)$. Defining $\Delta(t) := \mathbb{E}[L(\mathbf{G}(t+1)) | \mathbf{G}(t)] - L(\mathbf{G}(t))$ as the drift at slot t , the drift plus penalty (DPP) policy is designed to solve by observing only the current queue lengths $\mathbf{Q}(t)$ and the current network state $\boldsymbol{\omega}(t)$ on each slot t and then choosing $\alpha(t) \in A_{\boldsymbol{\omega}(t)}$ to minimize a bound on $\Delta(t) - V \sum_{u \in \mathcal{U}} \phi_u(\gamma_u(t))$. Here, $V > 0$ is a control parameter of the policy which affects a utility-backlog tradeoff. It is then easy to show that the resulting policy is given by the minimization, at each chunk time t , of the function:

$$\underbrace{k\mathbf{B}^T(t)\mathbf{Q}(t) - \mathbf{D}^T(t)\boldsymbol{\Theta}(t)}_{\text{congestion control}} - \underbrace{n\boldsymbol{\mu}^T(t)\mathbf{Q}(t)}_{\text{transmission scheduling}} - \underbrace{\left[V \sum_{u \in \mathcal{U}} \phi_u(\gamma_u(t)) - \boldsymbol{\gamma}^T(t)\boldsymbol{\Theta}(t) \right]}_{\text{obj. maximization}} \quad (13)$$

The choice of $m_u(t) \forall u \in \mathcal{U}$ affects only the term $k\mathbf{B}^T(t)\mathbf{Q}(t) - \mathbf{D}^T(t)\boldsymbol{\Theta}(t)$, the choice of $\boldsymbol{\mu}(t)$ affects only the term $n\boldsymbol{\mu}^T(t)\mathbf{Q}(t)$, and the choice of $\boldsymbol{\gamma}(t)$ affects only the term $V \sum_{u \in \mathcal{U}} \phi_u(\gamma_u(t)) - \boldsymbol{\gamma}^T(t)\boldsymbol{\Theta}(t)$. Thus, the overall minimization decomposes into three separate sub-problems. The first two sub-problems have a clear operational meaning,

and will be referred to as *congestion control* and *transmission scheduling*. Congestion control consists of choosing the quality index $m_u(t)$ for the chunk requested at time t and placed into the queue Q_u by every user u . Transmission scheduling consists of allocating the channel transmission rates $\mu_{hu}(t)$ for each helper h to its neighboring users $u \in \mathcal{N}(h)$. The third sub-problem involves the greedy maximization of each user network utility function with respect to the auxiliary control variables $\gamma_u(t)$.

A. Congestion Control: Pull Scheme

The congestion control sub-problem objective function (see (13)) can be explicitly expressed as

$$\sum_{u \in \mathcal{U}} \{kQ_u(t)B_{f_u}(m_u(t), t) - \Theta_u(t)D_{f_u}(m_u(t), t)\}.$$

The minimization of this quantity decomposes into separate minimizations for each user, namely, for each $u \in \mathcal{U}$, choose $m_u(t)$ equal to:

$$\arg \min_{m \in \{1, \dots, N_{f_u}\}} \{kQ_u(t)B_{f_u}(m, t) - \Theta_u(t)D_{f_u}(m, t)\}. \quad (14)$$

In order to implement this policy, it is sufficient that each user knows the length of its own request queue Q_u . Congestion control decisions are decentralized: each user u observes the length of its own *request queue* Q_u and requests the t -th chunk at a quality level according to (14). This policy is reminiscent of the current adaptive streaming technology for video on demand systems, referred to as DASH (Dynamic Adaptive Streaming over HTTP) [9], where the client (user) progressively fetches a video file by downloading successive chunks, and makes adaptive decisions on the source encoding quality based on its current knowledge of the congestion of the underlying server-client connection.

B. Transmission Scheduling

Transmission scheduling involves maximizing the weighted sum rate $\sum_{u \in \mathcal{U}} Q_u(t)\mu_u(t)$ where the weights are the request queue lengths (see (13)). Under our system assumptions for both PHY A and PHY B, this problem decouples into separate maximizations for each helper. Notice that here, unlike conventional cellular systems, we do not assign a fixed set of users to each helper. In contrast, the helper-user association is dynamic and changes as the users or helpers move around, and results from the transmission scheduling decision itself. We now describe the algorithm to solve the weighted sum rate maximization at each helper for the system PHY B.

Transmission Scheduling for PHY B: For each $h \in \mathcal{H}$, the transmission scheduling problem can be written as the following *Linear Program* (LP):

$$\text{maximize} \quad \sum_{u \in \mathcal{N}(h)} Q_u(t)\mu_{hu}(t) \quad (15)$$

$$\text{subject to} \quad \boldsymbol{\mu}_h(t) \in \text{coh}\{\mathbf{C}_h(\mathbf{S}_h, t) : \mathbf{S}_h \subseteq \mathcal{N}(h)\}. \quad (16)$$

where $\boldsymbol{\mu}_h(t)$ is the $|\mathcal{N}(h)|$ -dimensional vector of elements of the set $\{\mu_{hu}(t) : u \in \mathcal{N}(h)\}$. It is well known that the optimal value of an LP is attained at an extreme point/vertex

of the polytopal feasible rate region described in (16). It is also known that the convex hull of a set of points has a subset of those points as its set of vertices (Prop. 2.2 in [14]). Thus, the LP reduces to maximizing the weighted sum rate $\sum_{u \in \mathcal{N}(h)} Q_u(t)\mu_{hu}(t)$ over the discrete set of points $\{\mathbf{C}_h(\mathbf{S}_h, t) : \mathbf{S}_h \subseteq \mathcal{N}(h)\}$ of cardinality $2^{|\mathcal{N}(h)|} - 1$. This means that the subset which maximizes the weighted sum rate is chosen and the same subset is served throughout the entire chunk time slot t . This saves a lot of effort in terms of implementation (e.g., overheads in pilot dissemination in a TDD reciprocity-based system) when compared to approaches which switch between different active user subsets every OFDM block τ .

One can observe from (2) that when helper h schedules the subset \mathbf{S}_h of users for MU-MIMO beamforming, the rate of each user $u \in \mathbf{S}_h$ depends only on the cardinality S_h but not on the identity of the members of the subset \mathbf{S}_h . This implies that for a fixed subset size S , the subset of users maximizing the weighted sum rate can be obtained by sorting the users in $\mathcal{N}(h)$ according to the weighted rate $Q_u(t) \log \left(1 + \frac{g_{hu}(t)(M-S+1)P_h}{S(1+\sum_{h' \neq h} P_{h'u}g_{h'u}(t))} \right)$ and choosing greedily the best S users. This ‘sort+greedy algorithm’ is repeated for every subset size $S \in \{1, \dots, |\mathcal{N}(h)|\}$ and the subset which gives the maximum weighted sum rate is chosen. A typical sorting algorithm has complexity $O(|\mathcal{N}(h)| \log(|\mathcal{N}(h)|))$ and since the sorting is repeated for every subset size, our algorithm has complexity $O(|\mathcal{N}(h)|^2 \log(|\mathcal{N}(h)|))$ which improves upon existing user scheduling algorithms [15] for the MIMO broadcast channel. Note that the $n\mu_{hu}(t)$ video-encoded bits transmitted by helper h to user u should correspond to the chunks at the head of line of the request queue Q_u , assuming the quality chosen by user u in a previous time slot based on the pull scheme (14). Thus, each user u must also broadcast the metadata (chunk number and quality) of the chunks at the head of line along with Q_u to helpers in $\mathcal{N}(u)$. In addition, when a user is (possibly) served by multiple helpers in a single slot, the user gets the $n\mu_u(t) = \sum_{h \in \mathcal{N}(u)} n\mu_{hu}(t)$ video-encoded bits corresponding to the head of line chunks by progressively downloading different PHY layer sub-packets of a chunk (assuming a chunk is segmented into different PHY layer sub-packets for transmission) from different helpers. This would actually require the user to keep track of the PHY layer sub-packets that it has downloaded. However, if intra-session network coding is employed to encode the PHY sub-packets of each chunk, then the user no longer needs to bookkeep at the sub-packet level and just needs to download the required number of linear combinations to decode the chunk. Thus, our approach is ideally suited to an application of distributed storage codes like DRESS [16] which makes the sub-packet bookkeeping problem much simpler.

C. Greedy maximization of the network utility function

Each user $u \in \mathcal{U}$ keeps track of $\Theta_u(t)$ and chooses its virtual queue arrival $\gamma_u(t)$ in order to solve:

$$\text{maximize} \quad V\phi_u(\gamma_u(t)) - \Theta_u(t)\gamma_u(t) \quad (17)$$

$$\text{subject to} \quad D_u^{\min} \leq \gamma_u(t) \leq D_u^{\max}. \quad (18)$$

These decisions push the system to approach the maximum of the network utility function. By appropriately choosing the functions ϕ_u , we can impose some desired notion of fairness.

IV. ALGORITHM PERFORMANCE

Following in the footsteps of [1], it can be shown that the time average utility achieved by the policy comes within $O(\frac{1}{\sqrt{V}})$ of the utility of a genie-aided T-slot look ahead policy for any arbitrary sample path $\omega(t)$ with a $O(V)$ tradeoff in time averaged backlog. The proof is omitted due to space constraints.

V. PRE-BUFFERING AND RE-BUFFERING CHUNKS

The goal here is to determine the delay T_u after which user u should start playback, with respect to the time at which the first chunk is requested (beginning of the streaming session). We define the size of the playback buffer Ψ_t as the number of playable chunks in the buffer not yet played. Without loss of generality, assume again that the streaming session starts at $t = 1$. Then, Ψ_t is recursively given by the updating equation:¹

$$\Psi_t = \max \{ \Psi_{t-1} - 1 \{t > T_u\}, 0 \} + |a_t|. \quad (19)$$


where $|a_t|$ is the number of chunks which are completely downloaded in slot t . Let A_k denote the time slot in which chunk k arrives at the user and let W_k denote the delay with which chunk k is delivered. Note that the longest period during which Ψ_t is not incremented is given by the maximum delay to deliver chunks. Thus, each user u needs to adaptively estimate W_k in order to choose T_u . In the proposed method, at each time $t = 1, 2, \dots$, user u calculates the maximum observed delay E_t in a sliding window of size Δ , by letting:

$$E_t = \max \{ W_k : t - \Delta + 1 \leq A_k \leq t \}. \quad (20)$$


Finally, user u starts its playback when Ψ_t crosses the level ξE_t , i.e., $T_u = \min \{ t : \Psi_t \geq \xi E_t \}$. where ξ is a tuning parameter. If a stall event occurs at time t , i.e., $\Psi_t = 0$ for $t > T_u$, the algorithm enters a re-buffering phase in which the same algorithm presented above is employed again to determine the new instant $t + T_u + 1$ at which playback is restarted.

VI. NUMERICAL EXPERIMENT


We consider a simple topology with 2 helpers and 20 users as shown in Figure (1a). The parameters of the simulation are chosen to be identical to those chosen in the experiments in [1]. We choose the utility function $\Phi_u(\cdot) = \log(\cdot) \forall u \in \mathcal{U}$ to impose proportional fairness. We use the same video file in [1] with VBR encoded chunks and with video quality measured using the SSIM index [12]. We simulate the algorithms designed above and plot the CDF over the user population of video streaming QoE metrics. We observe that the average video quality performance is the same in all systems but there is a significant improvement of MU-MIMO PHY B over PHY A in terms of pre-buffering and re-buffering times. The blue curves correspond to system PHY A. The green curves correspond to PHY B with number of antennas $M = 10$ and maximum active user subset size $S_{\max} = 5$. The red curves correspond to PHY B with $M = 20$ and $S_{\max} = 10$.




(a)



(b) CDF of avg. video quality



(c) CDF of pre buffering time



(d) CDF of rebuffering percentage

REFERENCES

- [1] D. Bethanabhotla, G. Caire, and M. J. Neely, "Joint transmission scheduling and congestion control for adaptive video streaming in small-cell networks," *arXiv preprint arXiv:1304.8083*, 2013.
- [2] Cisco, "The zettabyte era-trends and analysis," 2013.
- [3] S. Sesia, I. Toufik, and M. Baker, *LTE: the Long Term Evolution-From theory to practice*. Wiley, 2009.
- [4] J. Hoydis, S. Ten Brink, and M. Debbah, "Massive MIMO: How many antennas do we need?" in *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2011, pp. 545–550.
- [5] J. Hoydis, M. Kobayashi, and M. Debbah, "Green small-cell networks," *Vehicular Technology Magazine, IEEE*, vol. 6, no. 1, pp. 37–43, 2011.
- [6] E. H. Ong, J. Knecht, O. Alanen, Z. Chang, T. Huovinen, and T. Nittala, "IEEE 802.11 ac: Enhancements for very high throughput WLANs," in *2011 IEEE 22nd International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2011, pp. 849–853.
- [7] Y. Yi and M. Chiang, "Stochastic network utility maximisation-a tribute to Kelly's paper published in this journal a decade ago," *European Transactions on Telecommunications*, vol. 19, no. 4, pp. 421–442, 2008.
- [8] M. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1–211, 2010.
- [9] Y. Sánchez, T. Schierl, C. Hellge, T. Wiegand, D. Hong, D. De Vleeschauwer, W. Van Leekwijck, and Y. Lelouedec, "iDASH: improved dynamic adaptive streaming over http using scalable video coding," in *ACM Multimedia Systems Conference (MMSys)*, 2011, pp. 23–25.
- [10] A. Begen, T. Akgul, and M. Baugher, "Watching video over the web: Part I: Streaming protocols," *Internet Computing, IEEE*, vol. 15, no. 2, pp. 54–63, 2011.
- [11] A. Ortega, "Variable bit-rate video coding," *Compressed Video over Networks*, pp. 343–382, 2000.
- [12] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *Image Processing, IEEE Transactions on*, vol. 13, no. 4, pp. 600–612, 2004.
- [13] H. Huh, A. M. Tulino, and G. Caire, "Network mimo with linear zero-forcing beamforming: Large system analysis, impact of channel estimation, and reduced-complexity scheduling," *Information Theory, IEEE Transactions on*, vol. 58, no. 5, pp. 2911–2934, 2012.
- [14] G. M. Ziegler, *Lectures on polytopes*. Springer, 1995, vol. 152.
- [15] T. Yoo and A. Goldsmith, "On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming," *Selected Areas in Communications, IEEE Journal on*, vol. 24, no. 3, pp. 528–541, 2006.
- [16] S. Pawar, N. Noorshams, S. El Rouayheb, and K. Ramchandran, "Dress codes for the storage cloud: Simple randomized constructions," in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*. IEEE, 2011, pp. 2338–2342.

¹ $1\{\mathcal{K}\}$ denotes the indicator function of a condition or event \mathcal{K} .